Direct Digital Synthesis

What is It and How can I Use It?

by Dr P H Saul, G8EUX

'DIRECT DIGITAL SYNTHESIS' (DDS) has appeared several times recently in *RadCom*, both in *Technical Topics* [1] and in a review [2]. Direct Frequency Synthesis (DFS) and Numerically Controlled Oscillator (NCO) are synonymous terms. It is the latest feature offered in top-of-the-range rigs such as the Icom IC-781 and the new Yaesu FT-1000, which has, according to the advertisements, no less than five direct digital synthesisers. What then is this apparently new technique, and what is its relevance to amateur radio?

TERMINOLOGY

FIRST, SOME TERMINOLOGY explained. Direct synthesis (without the digital bit) simply means the production of an output frequency, either a transmitter output or a local oscillator for a receiver by the addition of several, sometimes many oscillators, usually by a combination of mixing, multiplying, re-mixing etc., with many stages of filtering. A truly analogue technique, and horrifically expensive to carry out really well. Examples were primarily limited to specialist, often military, synthesizers in the fifties. Outputs could be made very clean, especially in respect of close-to-carrier phase noise, but only by massive amounts of filtering, usually racks full of it! Of course, synthesizers of this type never made it into amateur production equipment.

A simpler version did go into quantity production, primarily, at least in numbers, in the early U.S. market C.B radios. This was the 'crystal bank' synthesizer, shown in **Fig 1**. This arrangement used two arrays of crystals, selected in appropriate combinations to give the required coverage and channel spacing. It worked very well over a restricted frequency range, could offer very low phase noise since crystal oscillators are inherently 'quiet', and could be made acceptably (for the time) compact, and low power. A familiar example on the UK market some 15 years ago was the very well known 'Liner 2'. This used crystal bank synthesis, with a front-panel tuned VXO on a further crystal oscillator for the conversion up and down from the effective 'tuneable IF' at 30MHz to the working frequency of 145MHz. The need for the VXO illustrates the weakness of this type of synthesizer; close channel spacing is essentially prohibited by the very large number of crystals which would be needed. Triple bank versions were described in the professional literature, but were not produced for amateur purposes. The other disadvantage of this class of synthesizer is the high cost of the crystals, but it remains an interesting technique capable of very high performance with reasonable design. The well-documented 'faults' of the Liner 2 were not connected with the synthesizer, but were due to overdrive of a later mixer stage.

Most of the earlier commercial amateur bands HF radios overcame the need for frequency stability by the use of a tuneable IF system and a carefully engineered VFO. Examples are in the KW range, the Yaesu/Sommerkamp FT-DX series and, in homebrew form, the G2DAF. This technique survived in solid state form, with among others, the FT101. However, in parallel with these came the early attempts at combining the best VFO characteristics with synthesis in order

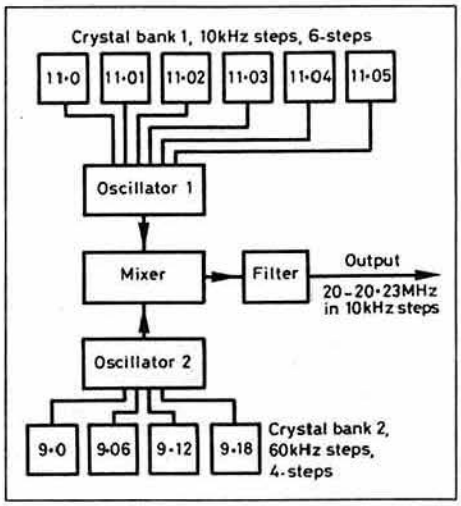

Fig 1: Simplified crystal bank synthesizer

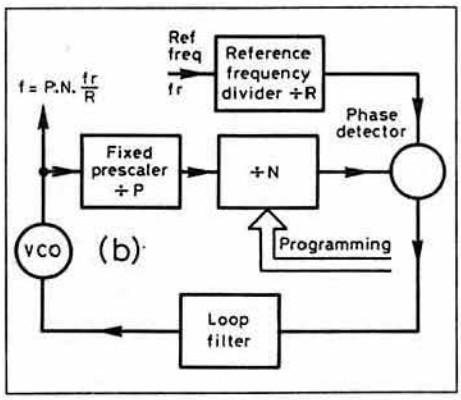

Fig 2a:Direct division

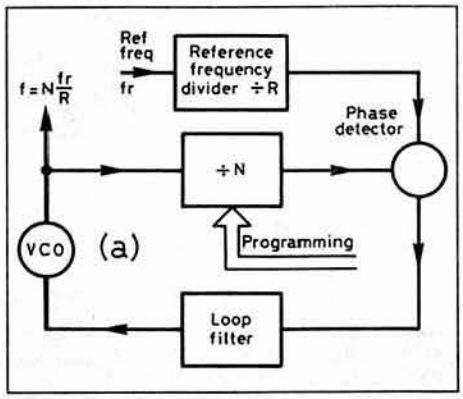

Fig 2b: Fixed prescaling

to build a single conversion receiver, with the objective of achieving exceptionally wide dynamic range. An example of this was G3PDM's classic phase-locked oscillator as used in his receiver [3], This was a true synthesizer, albeit analogue, based primarily on valves, with a VFO using a FET and two bipolar transistors which is still regarded as the standard to beat. The author was present at a demonstration of this receiver where a signal below 1 microvolt suffered no apparent degradation from a 10volt signal 50KHz away, an amazing 140+dB dynamic range made possible by the extremely low phase noise of this synthesizer, and of course the superb linearity of a beam-deflection mixer run with 100 volt push-pull local oscillator drive!

About this time (1969), the first truly digital (but not direct) synthesizers were beginning to appear, chiefly in military equipment; the outstanding example is the UK sourced 'Clansman' series of radios, still in service with UK forces.

BASIC SYNTHESIZER

A BASIC SINGLE LOOP digital synthesizer is shown in **fig 2**. Actually, this is very simplified. Fully programmable dividers of appropriate frequency range were not immediately available. Two solutions were proposed for this; the 'dual modulus' prescalers and a variant of the crystal mix scheme. The first of these uses the properties of a dual modulus, or in some cases even a triple modulus, counter to overcome the difficulty of producing a really high frequency fully programmable divider. The scheme is shown in **fig 3**. It consists of a VCO, usually in discrete component form, a dual modulus counter, a main counter, a reference oscillator and divider chain, a phase detector and a loop amplifier/filter. This type of synthesizer has been described many times, both in the professional and amateur fields. In some variants, it probably represents the current status of PLL frequency synthesis, especially for very wide frequency range applications such as TV receivers and satellite tuners.

The second method, that of crystal mixing, is probably the most widely used technique in amateur and PMR equipment, where wide operating frequency range is not needed. It offers advantages of simplicity and low component count, especially when the digital functions of the system can be contained in a single chip. The scheme is shown in **fig 4**. Unlike the dual modulus scheme, no prescaler is needed, but a high frequency mixer is necessary, usually with two crystals , one for mixing down and one for the reference.

Almost all the commercially available phase locked loop synthesizers use an external VCO, which must be designed by the equipment builder; not too difficult to a commercial concern, but no mean task for the amateur building a one-off. The problem is that the VCO is critical to achieving good performance from the synthesizer,

especially in respect of phase noise , and hence reciprocal mixing in a receiver. While the lock loop can take care of phase noise close to the carrier, a badly designed circuit or a badly laid out PCB can cause a synthesizer to 'feature' sidebands on the output spectrum at the comparison frequency and sometimes harmonics of it. The VCO dominates the noise away from the nominal frequency, ie outside the PLL bandwidth, which may be only a kHz or so, less in the case of very narrow frequency increments. Phase detector design is another critical area; the most recent PLL chips [5] offer two or more phase detectors to cover the far-from-lock and locked-in cases. All in all, PLL synthesizer design is a very demanding and essentially analogue task, as anyone who has built one and carefully examined the output on a spectrum analyser will testify. Nevertheless, truly excellent performance can be achieved, by the professional engineer in factory built equipment and by the very dedicated and well equipped amateur who does not cost his time into the project. It should be said that for any synthesizer work, a spectrum analyser, or at least a continuously tuneable receiver covering the frequency of interest and a harmonic or two, is essential.

LOCK-UP TIME

PHASE LOCKED LOOP synthesizers have one further disadvantage. The lock-up time is inherently related to the smallest frequency increment. For an FM only rig, with say 5KHz minimum frequency step, this is of no real consequence. For SSB, however, even 100Hz steps are disconcertingly large to most operators, and 10 Hz or better is preferred for the 'analogue feel'. Professional systems overcome this by complex, multiple loop synthesis, which can be both expensive, and when the filters are taken into account, bulky.

Amateur equipment practice has tended to take a more pragmatic route. Minimum step size in the synthesizer is usually 1KHz, while increments between are achieved with a separate control knob on older equipment, or by a digital analogue converter (DAC), operated by the last digit of the frequency set. The analogue voltage tunes a VXO in the rig, which may either be the conversion crystal in the synthesizer, or may be the reference crystal itself. Care must be taken to ensure that the pulling range is accurate or the '100Hz' steps will show a jump in one direction or the other at the 1KHz increments. Actually, many rigs do show this if observed carefully; the reason is that, until recently the DAC was a fairly simple affair consisting of a resistor array and switching transistors. More recently, real DACs have been used, but this is not a complete solution, since the VXO is unlikely to be linear to the required degree, so some step nonlinearity is inevitable. At least if the end points are not seriously wrong, this should not be a problem. Some rigs do tune in 10Hz steps, especially on HF; this is an extension of the technique to 100 steps instead of just 10. Again, there is the issue of the 1KHz crossover points; but with 10Hz steps this can be made less noticeable on a well designed and adjusted rig. This is the final point on PLL synthesis; properly designed, it works very well indeed, but who has not heard of 'synthesizer whine', or read reviews of otherwise excellent radio limited in performance by synthesizer noise causing reciprocal mixing? My own experience includes a rig which preferred not to go mobile; the synthesizer lock loop would go out of tune after a surprisingly repeatable mileage, only finally cured by the application of much wax to the VCO coil.

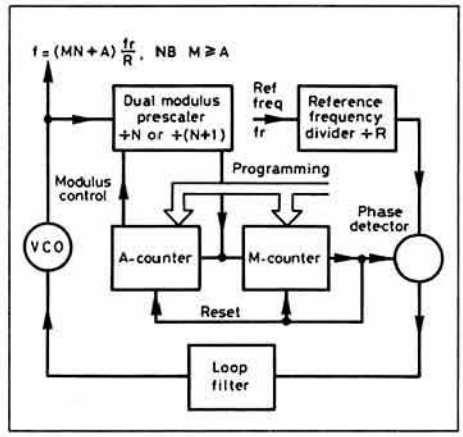

Fig 3: Dual Modulus prescaling

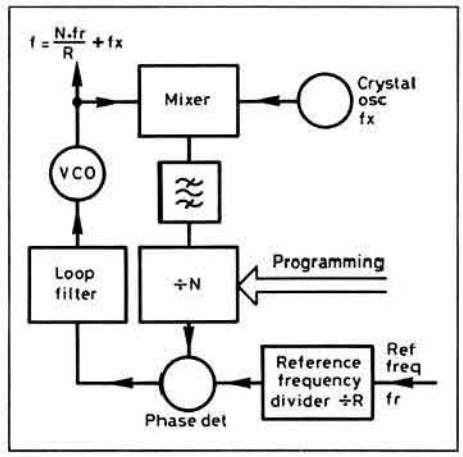

Fig 4: Mixing in the loop

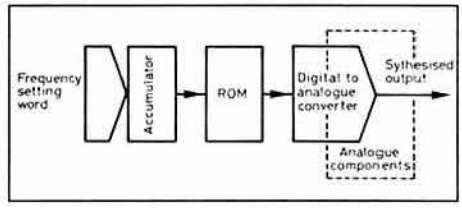

Fig 5: Direct frequency synthesizer

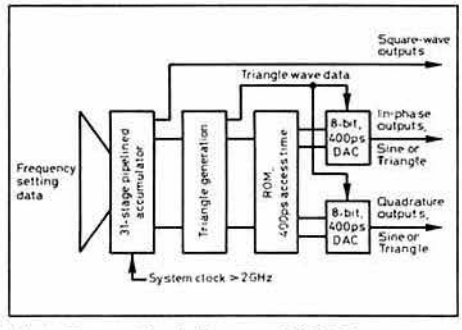

Fig 6: System block diagram of SP2002

A CURE-ALL?

WILL DIRECT DIGITAL synthesis cure all these ills? The answer is "no, not yet", but in the fairly near future this may turn into a qualified "yes". The basic direct digital synthesizer consists of an arrangement to generate the output frequency directly from the clock and the input data. The simplest conception is shown in **fig 5**. This consists of a digital accumulator, a ROM containing the pattern in digital form of a sinewave, and a digital to analogue converter. In general, the DAC will be followed by some filtering usually a low-pass filter somewhere just below half the clock frequency.

Dealing with the accumulator first, this is simply an adder with a store at each bit. It adds the input data word to that in the store. The input data word only changes when the required frequency is to be changed. In the simplest case the length of accumulator is the clock frequency divided by the channel spacing, although it is usually calculated the other way round, ie if a 5kHz channel spacing up to 150MHz is needed, a clock frequency of at least 300MHz will be required, due to the Nyguist sampling theorem, and 300MHz divided by 5KHz is 60,000. More conveniently, a 16 bit accumulator gives a 65,536 steps. If the step size is to be 5KHz, then a clock frequency to the accumulator of 65,536 x 5KHz is needed, i.e. 327.68MHz.

These numbers are the basic values chosen for the first commercially available VHF DDS, the Plessey SP2001. What the numbers mean is that by applying a binary number multiplied by 5KHz, ie the number to be applied is the 'channel number'. This makes programming particularly easy, especially as the device has a parallel input data format: DIP switches were used on the evaluation boards. The device produces any frequency in the range covered, ie 5KHz to over 100MHz in a single range, without any tuned circuits, although of course it does depend for stability on the clock source, which would normally be crystal controlled. Other frequency increments are available; any multiple of 5kHz by selection of input data, and others by choice of clock frequency, eg 6.25kHz requires a clock at 204.8MHz, with a two channel (2 x 3.125kHz) program word. Of course, frequency multiplication or mixing can be used to take the output to higher frequencies.

NEW DEVICE

RECENTLY, A NEW DEVICE (SP2002) has been announced [6]. This device, which is shown in block diagram form in **fig 6**, is more complex than the SP2001, since it has on chip DACs and facilities for square, triangle and sine outputs. Two DACs are needed because both phase and quadrature signals are available in true and complement form. The most significant bit (MSB) from the accumulator feeds the square wave output buffer direct. In parallel, the next 7 bits from the accumulator, which digitally represent a sawtooth waveform, feed a set of XOR gates, under control of the MSB, so that a triangle output is generated, in digital form, at this point in the circuit. Actually, two triangles in quadrature are generated. These can be digitally steered to the output DACs, or can be used to address a ROM containing data for sine and cosine waves; only 90 degrees is needed, since all four quadrants can be generated from one. Finally, if selected, the digital sine/cosine is fed to the DACs for conversion to analogue form.

This device was designed to operate at up to 500MHz output frequency, with 1Hz steps, so the full range is 1Hz to 500MHz, again in a single 'range' with no means of or requirement for tuning. The clock frequency is necessarily very high; the quadrature requirement adds a further factor 2, so nominal clock frequency is 2^31Hz, i.e. 2.147483648Ghz. This must be supplied with crystal controlled stability, although since phase noise is effectively divided down in the synthesizer, and the frequency is fixed, it can be generated by a simple multiplier from a crystal source. Output spectra are shown in **figs 7 and 8**. Fig 7 shows a clean output at exactly one quarter of the clock frequency. Noise floor is essentially generated by the spectrum analyser. Fig 8, however, is more representative of the general case; it shows a spectrum a 225MHz, a

frequency not integrally related to the clock frequency. Here the spurs have come up to a level about 50db below the carrier. This is the fundamental limitation of the DDS technique, and where most development work is going in the future. The limit comes from the finite word size and accuracy of the DAC, and incidentally the ROM, although this could have been made bigger fairly easily.

Fast DACs are difficult to make accurately, for two reasons. The first is technological; IC processes in general limit to a component matching accuracy of about 0.1%, ie 9 or 10 bits accuracy in a careful design. Other techniques, such as laser trimming of the resistors, can be used on some processes. However this tends to use older, slower processes, and especially requires large resistors for trimming, which slow the DAC settling time. The second problem is simply the requirements on fast settling; the DAC is required to get to the final value quickly, and this will in general happen with the smallest number of bits. An eight bit system was chosen to fit the process capabilities and the device requirements, but this leads to a limitation in the high level of spurious signals present. Basically, although some frequencies are very clean, in the worst case the spurious level is 6N dB below the carrier, where N is the number of effective DAC bits. For an 8 bit system, this give -48dB. Oversampling, ie running the clock at more than twice the output frequency gives some improvement, at 3dB per octave, by improvement of the DAC resolution, but only up to a limit of the DAC's accuracy. Typically this is about 9 bits or -54dB. In some applications, this may not be serious, since filtering can remove all but the close in spurs; phase locked translation loops into the microwave region also act as filters of relatively narrow band, while retaining the 1 Hz step capability.

In the amateur rigs available using DDS, the synthesizers operate at relatively low frequencies and are raised to the working frequency by PLL techniques; complicated, but still making the fine frequency increments available without compromise of PLL design. The devices used are probably CMOS types, with DAC accuracy around 12 bits. This gives spurious signals which should theoretically be -72dB referred to the carrier, which is adequate in most cases especially with some filtering from the PLL.

What of the homebrewer? Where does he get the devices? Does he want them anyway? All the devices mentioned by part number are commercially available; although the exact types in the Japanese equipment are not known to the author, US types of similar performance are also available. However, prices are likely to be high, at least for the next year or two. This is especially true for the VHF and UHF types, which are not really ideal anyway for single frequency (as opposed to 'hopping') radio receiver applications.

PRACTICAL DDS

A VERY USEFUL ARTICLE on a practical DDS appeared some years ago [7]. This showed how a synthesizer could be made using available TTL and a ROM. However, the output frequency was limited to 1MHz. Real DDS systems could be built by the amateur, and I intend to explore this for my own use. This is not intended to be a constructional article, so the comments made here are more in the nature of hints, without guarantee of success, but they outline a possible route to follow.

First, what do we need? An output up to say 10MHz would be useful for further conversion upwards in a mixer or PLL. Steps of 10Hz or

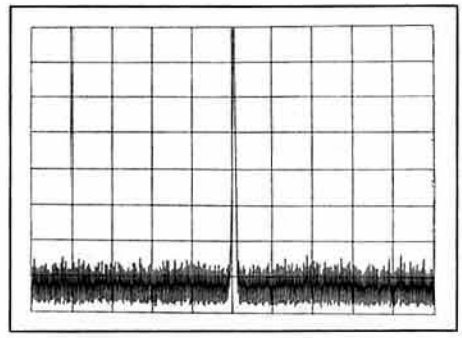

Fig 7: Output spectrum at 250MHz, with a 1GHz clock. 10MHz/div, X axis; 10dB/div, Y axis.

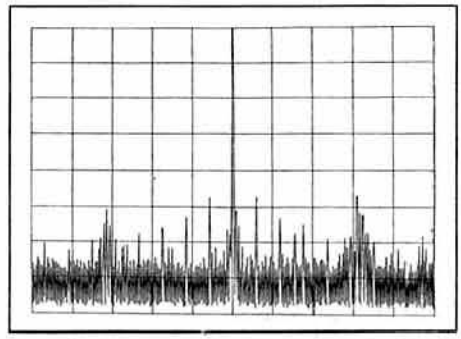

Fig 8: Output spectrum at 225MHz, with a 1GHz clock. 10MHz/div, X axis; 10dB/div, Y axis.

better are essential. This determines the clock frequency (>20MHz) and the accumulator size (21 bits). For convenience, the clock should therefore be 20.971520MHz. Standard TTL will work at this frequency although I would prefer HC CMOS.

A square wave would fit most applications in radio; there is no real need for a sinewave in local oscillators. However, the MSB output mentioned above has one problem; even higher spurs! Think of it as a one-bit DAC. Then -6NdB is not very good spurious suppression. The reason is that the square wave can give only increments in time, limited to the clock period in resolution. The sinewave also interpolates in amplitude, to the degree offered by the DAC accuracy, and so gives much better spurious control. However, if we are prepared to accept harmonics, which can be filtered out later if required, then a triangle is just as good in this respect as a sinewave. Since most commercial ROMs have access times of around 100ns, 10MHz is the limiting update rate, so a triangle has the further advantage of not needing the ROM at all. The DAC remains the major problem. Current low cost DACs are limited to 8 bits and settling times of over 1 microsecond.

A BETTER APPROACH

A MUCH BETTER APPROACH would be to use a current output DAC such as the SP9768 (8 bits, 5ns) or SP9770 (10bits, 12ns). These are ECL input devices, so will need interfaces to the CMOS accumulator and XOR gates; an ECL accumulator and XOR array should be considered. The latter device would give some 60db spurious suppression alone. Alternatively, it should be possible to add more bits using external discrete components, probably one further LSB and possibly two MSBs. Although temperature tracking of these components would not be good, the system would be operating in a fairly benign environment, so this is not critical. The DAC so produced would therefore be 12 or 13 bits accurate, the accuracy of the added bits coming from preset 10 turn potentiometers and some careful DC measurements. A full 4 digit DVM (or better) will be required for this operation. Of course, a better solution would be a fast 12 or more bit DAC; these are becoming available, but at a price.

Putting the system together should be fairly simple, using good VHF practice in order to avoid accidental cross talk onto the output lines. Power supplies need good decoupling for the same reason. A clock oscillator at this frequency should not pose problems. The only set up adjustment is the DAC if built up as above; ideally, DDS systems are set up free. Finally a check on a spectrum analyser will show how well it works.

CONCLUSION

IN CONCLUSION, DIRECT frequency synthesis is a promising technique for the future. There are certainly obstacles to overcome before it is generally used without a phase locked loop for filtering and conversion in HF radio applications, but it is already in use in amateur equipment. Eventually, we may see faster, more accurate DACs and new filtering techniques combine to produce a 100dB clean synthesized oscillator, with no set-up procedures at all and with extremely narrow channel spacing at low cost.

REFERENCES

1) Pat Hawker, 'Technical Topics' *RadCom*, December 1988 pp 957 - 958
2) Peter Hart 'ICOM IC-725 HF Transceiver' *RadCom* September 1989 pp 56 - 58
3) Peter Martin 'A receiver with noise immunity and Frequency Synthesis', *Radio Communication Handbook* (RSGB), fourth Edition pp 10.104 - 10.108.
(4) Plessey Semiconductors 'Frequency Dividers And Synthesizers IC Handbook', especially pp 304 - 306.
5) Plessey Semiconductors Data Sheet SP8853
6) Plessey Semiconductors Data Sheet SP2002
7) JHJ Dawson 'Direct Digital Frequency Synthesizer' Wireless World December 1981 pp 40 - 43.

CORRECTION

Avoiding Wind Vibration Damage to Aerials by A Bolton, GM3BMI.

In keying in the mathematics for Mr Bolton's article (Sept *RadCom*), the following errors crept into the Appendix:

Page 48, section 2: In 'Fatigue Stress', line 7, replace '1' with 'l'.

Page 48, section 3: In 'Natural frequencies', line 9, the '(4)' is a reference, not part of the equation.

Page 48, col 2, 3 lines from the bottom: After 'rotate', insert 'the upper with both ends pinned but free to rotate'.

Page 48, col 3, section 4, equation (6): Replace '1' with 'l'.

Page 49, col 1, 'c:natural frequency' should, of course have been in larger bold letters and not indented.

We hope that readers did not find that these errors detracted from the usefulness of the article.